A BAYESIAN OPTIMISATION APPROACH FOR MODEL INVERSION OF HYPERSPECTRAL-MULTIDIRECTIONAL OBSERVATIONS: THE BALANCE WITH *A PRIORI* INFORMATION

Wout Verhoef [a,b]

[a] National Aerospace Laboratory NLR, P.O. Box 153, 8300 AD Emmeloord, The Netherlands
[b] International Institute for Geo-Information Science and Earth Observation ITC, 7500 AA Enschede, The Netherlands
verhoef@nlr.nl, verhoef@itc.nl

KEY WORDS: model inversion, hyperspectral, multidirectional, SAIL, MODTRAN, PROSPECT, Bayesian optimisation

ABSTRACT:

Hyperspectral-multidirectional radiance observations of the land surface from space potentially form one of the richest sources of geobiophysical information possible. For soil-vegetation objects, the retrieval of this information can be simulated by radiative transfer modelling. In combination with a couple of atmospheric parameters, the surface reflectance model SLC (soil-leaf-canopy) has more than twenty degrees of freedom, which all have a potential impact on top-of-atmosphere radiance data in hyperspectral-multiangular feature space. With such a high dimensionality, model inversion methods like look-up table techniques and neural networks tend to become less practicable, and cost-function optimisation re-emerges as a viable alternative. However, model inversion by optimisation techniques is often plagued by numerical instability due to the so-called ill-posedness of the model inversion problem. In the present paper, this ill-posedness of the problem is investigated and diagnosed by means of a singular value decomposition (SVD) of the Jacobian matrix, which contains the partial derivatives of all observations with respect to the model variables. In addition, it is demonstrated how in a Bayesian approach the incorporation of *a priori* information can increase the numerical stability of the model inversion. This leads to an extremely efficient optimisation algorithm, which for randomly selected model variable data reaches an adequate solution in about 99% of the cases, in less than twenty iteration steps. The paper will introduce the model SLC, its coupling with the atmosphere, for which MODTRAN4 is used, and for some selected cases it will analyse the SVD results in order to explain the causes of ill-posedness. A few model inversion sequences will be presented in order to illustrate the numerical stability of the algorithm and its ability to reach a plausible solution under various circumstances. The speed of this method is still limited, but it might be applied selectively to representative pixels in a field, or to "calibrate" the fixed model parameters in a low-dimensional look-up table or neural network model inversion solution.

1. INTRODUCTION

Traditionally, remote sensing satellite missions for earth observation over land have often been designed to provide a range of geobiophysical products to a wide and globally distributed community of users. Operational products based on this philosophy are for instance MODIS LAI and MODIS fAPAR (Knyazikhin et al., 1998), which employ knowledge obtained from vegetation-soil radiative transfer modelling and global maps of ecotypes, but nevertheless are still basically rather sensor-specific. However, there is a growing awareness that algorithms designed for the massive processing of earth observation data into single products from single satellite missions might interfere with the consistent use of several land surface variables in dynamic process models. Also, different providers might disseminate basically the same product (e.g. LAI) derived from different sensors and based on different algorithms, making it hard for a user to decide which LAI is the most suitable for the intended application. Therefore, it becomes more and more obvious that a multisensor / multimission approach might be more successful in providing the user community with a range of products that can be assimilated in local process models in a more natural and self-consistent manner. In this way also the quality of products could be improved, since more information from various sources, including local information stored in GIS, could be integrated, thus reducing the chance of inconsistencies. In Figure 1 it is illustrated how such an approach could be applied to assimilate data from one satellite sensor. The surface object properties are stored in a GIS and by means of a generic remote sensing (RS) earth observation model which includes the atmosphere, top-of-atmosphere (TOA) radiance image data are simulated that have been adapted to the spatial, spectral and geometric properties of the sensor that produced the actual image. Comparison of both images leads to conclusions on the adjustment of object properties and these are then fed back until actual and simulated images sufficiently match. This feedback loop establishes model inversion on the level of complete (series of) images by forward simulation, with the great advantage that the impact of the complexity of the heterogeneous landscape and the atmosphere on image formation, including topography, can be captured fairly

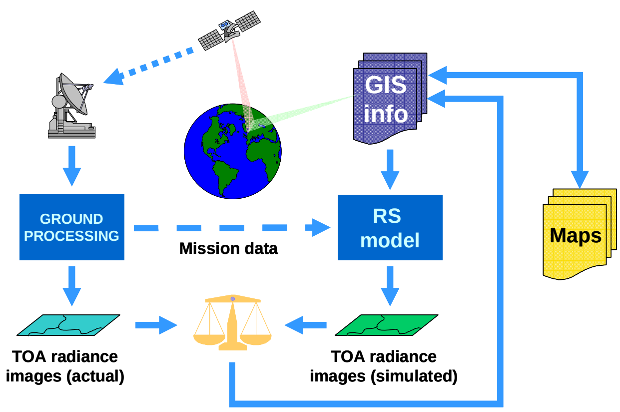

Fig. 1 Updating geographic information by balancing new earth observation data with existing forward-modelled *a priori* information.

well by forward modelling, making various atmospheric and topographic corrections unnecessary. Besides, since complete images are compared, the correct values of atmospheric parameters will be found in an early stage, as these have an effect on all pixels in the image, and thus the overall cost function for the whole image, which is to be minimised during the optimisation process, will rapidly decrease if a better characterisation of the atmosphere is found. This concept of parameter retrieval by forward modelling can be characterised by the following features:

- Simultaneous retrieval of all object properties
- Self-consistent set of object properties is obtained
- High degree of sensor-independence
- User-oriented because of local *a priori* GIS info

If current surface properties stored in the GIS are to be updated using a new satellite image, of course the question arises what should be the resistance against modifying data. Here the balance between new earth observation data and existing GIS info comes into play, and a key element one could use to base decisions on is uncertainty. If there is little uncertainty about the correctness of the current GIS info, then the resistance against changing it must be high, even if satellite earth observation data seem to indicate that there are discrepancies. On the other hand, if one actually does not know whether the current info is correct, one will be inclined to readily accept anything that can be derived from a new satellite observation. The question addressed in this paper is how one can balance both sources of information in an optimum way. For this, the model inversion of hyperspectral multidirectional radiance observations from space was taken as a prototype, since this is an ultimate example of the richness of earth observation data, for which the inversion is a challenge because of the complexity of the data, the high dimensionality and the chance of numerical problems due to ill-posedness of the model inversion problem. In what follows, first the set of coupled models used is introduced in section 2. Next the theory of Bayesian model inversion is discussed, and in section 4 some results of model inversion are presented.

2. MODELLING SET-UP

The set of models used was meant to represent the generation of TOA hyperspectral radiance data for generic soil-vegetation objects. The integrated soil-leaf-canopy model SLC (Verhoef & Bach, 2007) was used to generate surface reflectances for the former candidate mission SPECTRA over the wavelength range 400-2400 nm at 10 nm resolution and under 7 directions, as representative of data-rich inputs. The SLC model consists of the following submodels:

- 4-stream modified Hapke (1981) BRDF model
- Soil moisture effect after Bach & Mauser (1994)
- PROSPECT (Jacquemoud & Baret, 1990) leaf model
- 4SAIL2 canopy RT model (including canopy - soil interaction)

Compared to previous versions of SAIL, advancements in the 4SAIL2 model (Verhoef & Bach, 2007) can be summarized as follows:

- leaf colours different in two layers
- crown clumping effect included
- output of spectral canopy absorption (support for fAPAR) and observed fractional vegetation cover (FVC)
- numerically robust (singularities intercepted)
- speed-optimised

A complete list of object properties of the SLC model is presented in Fig. 2

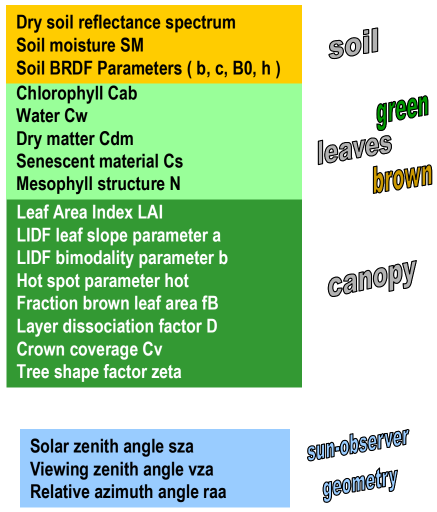

Fig. 2 Object properties of SLC model

For a given dry soil reflectance spectrum the number of free parameters is 5 for the soil, 5 for green leaves, 5 for brown leaves, and 8 for canopy structure, so 23 in total. In order to simulate observations from space of top-of-atmosphere radiances, MODTRAN4 was coupled to SLC as illustrated in Fig. 3. This coupling allows also to simulate products like fAPAR and surface albedo, but this falls outside the scope of the present paper. As an illustration, Fig. 4 shows a sample of simulated TOA observations of hyperspectral radiances under 7 directions for a soil-vegetation object with an LAI of one under standard conditions in April at mid-latitudes for an atmospheric visibility of 23 km.

For model inversion, the sensitivity of TOA radiance observations to changes in surface variables is particularly important. However, for successful inversions this is not sufficient. If two variables induce similar changes in TOA radiance spectra, it will be hard to identify which variable caused the change, so changes caused by one variable should also be linearly independent from changes caused by other variables.

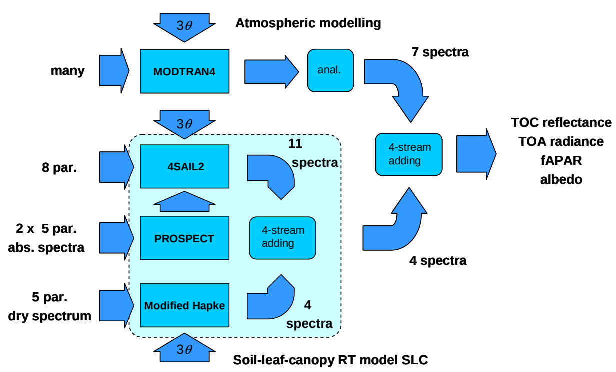

Fig. 3 Coupling of SLC outputs to MODTRAN4 for simulation of TOA radiance spectra.

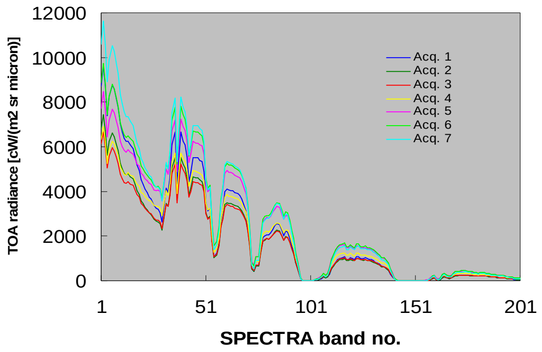

Fig. 4 Sample of simulated hyperspectral-multidirectional TOA radiance observations for a soil-vegetation object

Figures 5 and 6 show examples of the sensitivity of spectral and directional TOA radiances to a selection of surface variables. These illustrate that the sensitivities to the various parameters are spectrally and angularly quite diverse, although one can also notice similarities. Especially for the canopy structure the shapes of spectral sensitivities or often similar. Here, different parameters often give different angular responses, so that multiangular information provides quite some discriminative power. The amount of linear independence in the sensitivity to surface variables can be established by investigating the Jacobian matrix, which contains the partial derivatives of all observables (spectral bands and viewing directions) with respect to all surface variables. A very useful tool for this is singular value decomposition (SVD). Any matrix can be decomposed by SVD, and for a Jacobian matrix $\boldsymbol{J}$ one obtains

$$\boldsymbol{J} = \boldsymbol{U}\boldsymbol{S}\boldsymbol{V}^{\mathrm{T}} , \qquad (1)$$

with $\boldsymbol{U}^{\mathrm{T}}\boldsymbol{U} = \boldsymbol{I}$; $\boldsymbol{V}^{\mathrm{T}}\boldsymbol{V} = \boldsymbol{V}\boldsymbol{V}^{\mathrm{T}} = \boldsymbol{I}$; $\boldsymbol{S}$ diagonal.

The sensitivity of model observables to surface parameters is thus described by

$$\Delta\boldsymbol{r} = \boldsymbol{J}\Delta\boldsymbol{p} = \boldsymbol{U}\boldsymbol{S}\boldsymbol{V}^{\mathrm{T}}\Delta\boldsymbol{p} ,$$

which after pre-multiplication by the transposed of $\boldsymbol{U}$ gives

$$\boldsymbol{U}^{\mathrm{T}}\Delta\boldsymbol{r} = \boldsymbol{S}\,\boldsymbol{V}^{\mathrm{T}}\Delta\boldsymbol{p} . \qquad (2)$$

This result expresses that it is possible to obtain a vector of linearly transformed changes of surface variables, which are one-to-one related to the elements of a vector of linearly transformed changes of observables. The relation connecting both transformed vectors is given by the diagonal matrix of singular values $\boldsymbol{S}$. One may state that the singular values express the local sensitivities of a set of linearly transformed observables to a set of linearly transformed surface variables. Linear dependence will be expressed by one or more singular values being equal to zero, indicating that there are certain linear combinations of observables which have no sensitivity to a space of surface variables (the so-called null-space). This is exactly the situation encountered when model inversion becomes an ill-posed problem. In that case the Jacobian matrix is singular, so that the model inversion problem has no solution, or rather has multiple solutions. Thus, singular value decomposition is a useful tool to diagnose ill-posedness. In the next section it will be shown how the ill-posedness can be reduced by using *a priori* information.

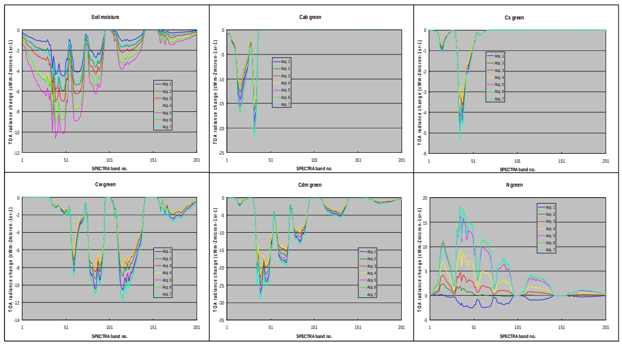

Fig. 5 Sensitivities of hyperspectral-multidirectional TOA radiance observations to soil moisture and leaf optical properties. Top row: soil moisture, leaf chlorophyll, brown pigment; bottom row: leaf water, dry matter, *N* parameter.

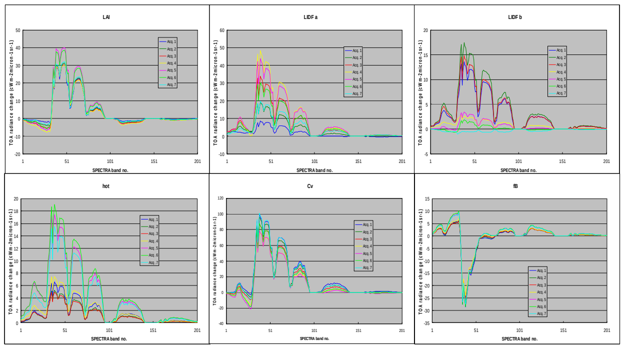

Fig. 6 Sensitivities of hyperspectral-multidirectional TOA radiance observations to canopy structure. Top row: LAI, LIDF a parameter, LIDF b parameter; bottom row: hot spot parameter, crown cover, fraction brown leaves.

3. BAYESIAN MODEL INVERSION

If Newton's optimisation method is applied without any regularisation by means of *a priori* information, in the neighbourhood of the final solution one can write

$$\Delta\boldsymbol{r} = \boldsymbol{J}\Delta\boldsymbol{p} ,$$

where $\Delta\boldsymbol{p}$ is the change in the surface parameters required to remove the discrepancy $\Delta\boldsymbol{r}$ between measured and modelled observables. Formally, the solution in terms of a change of parameters which removes the discrepancy in the observables is found by pre-multiplication with the transposed of $\boldsymbol{J}$, giving

$$\boldsymbol{J}^{\mathrm{T}}\Delta\boldsymbol{r} = \boldsymbol{J}^{\mathrm{T}}\boldsymbol{J}\Delta\boldsymbol{p} \text{ , or } \Delta\boldsymbol{p} = (\boldsymbol{J}^{\mathrm{T}}\boldsymbol{J})^{-1}\boldsymbol{J}^{\mathrm{T}}\Delta\boldsymbol{r}$$

However, this solution fails if the Jacobian matrix is singular, since one can show that the matrix

$$(\boldsymbol{J}^{\mathrm{T}}\boldsymbol{J})^{-1}\boldsymbol{J}^{\mathrm{T}} = \boldsymbol{V}\boldsymbol{S}^{-1}\boldsymbol{U}^{\mathrm{T}} ,$$

which obviously leads to infinite changes in parameters if any of the singular values equals zero.

Regularisation of Newton's iteration method can be accomplished by mixing it with the solution which goes into the direction of the *a priori* parameter vector, and using proper weights.

For the optimisation of a single parameter p, for which a model provides the solution p_m, and the *a priori* value is p_a, the weights are related to the uncertainties attached to the model result and the *a priori* value. If these are expressed by their variances $\sigma_m^{\ 2}$ and $\sigma_a^{\ 2}$, the Bayesian final solution is given by

$$p = \frac{\sigma_a^{\ 2} p_m + \sigma_m^{\ 2} p_a}{\sigma_a^{\ 2} + \sigma_m^{\ 2}} .$$

This can also be written as an equivalent expression which balances the differences with the model solution and the *a priori* solution, where the inverse variances are used as the weights:

$$\frac{p - p_m}{\sigma_m^{\ 2}} = -\frac{p - p_a}{\sigma_a^{\ 2}}$$

For normalised parameters having unit *a priori* variance, one can write

$$S^2 (p - p_m) = -(p - p_a) ,$$

where S is the relative sensitivity of the model to changes in normalised parameters. In this relative sensitivity the noise level of the sensor should be taken into account: the higher the noise level, the lower the relative sensitivity should be. For a multivariate system one can write similarly for a transformed variable space in the equilibrium situation

$$\boldsymbol{S}^2 \boldsymbol{V}^{\mathrm{T}} (\boldsymbol{p} - \boldsymbol{p}_m) = -\boldsymbol{V}^{\mathrm{T}} (\boldsymbol{p} - \boldsymbol{p}_a) .$$

The goal of an iteration step in the regularised Newton optimisation method is trying to minimise the difference with the equilibrium solution, which is given by

$$\boldsymbol{p}_s = \boldsymbol{V} (\boldsymbol{S}^2 + \boldsymbol{I})^{-1} (\boldsymbol{S}^2 \boldsymbol{V}^{\mathrm{T}} \boldsymbol{p}_m + \boldsymbol{V}^{\mathrm{T}} \boldsymbol{p}_a) , \tag{3}$$

and which clearly illustrates that more weight is given to the model solution if relative sensitivity is high, whereas in the case of no sensitivity at all the *a priori* solution is taken.

For the current vector of parameters one can write

$$\boldsymbol{p} = \boldsymbol{V} (\boldsymbol{S}^2 + \boldsymbol{I})^{-1} (\boldsymbol{S}^2 \boldsymbol{V}^{\mathrm{T}} \boldsymbol{p} + \boldsymbol{V}^{\mathrm{T}} \boldsymbol{p}) ,$$

and by subtracting this from Eq. (3) one obtains an updating rule which can be used in a regularised Newton optimisation algorithm, and reads

$$\Delta \boldsymbol{p} = \boldsymbol{V} (\boldsymbol{S}^2 + \boldsymbol{I})^{-1} [\boldsymbol{S}^2 \boldsymbol{V}^{\mathrm{T}} (\boldsymbol{p}_m - \boldsymbol{p}) + \boldsymbol{V}^{\mathrm{T}} (\boldsymbol{p}_a - \boldsymbol{p})]$$

Since in the neighbourhood of the solution the model can be linearised as given by Eq. (2), one can write

$$\boldsymbol{V}^{\mathrm{T}} (\boldsymbol{p}_m - \boldsymbol{p}) = \boldsymbol{S}^{-1} \boldsymbol{U}^{\mathrm{T}} (\boldsymbol{r}_o - \boldsymbol{r}) ,$$

where $\boldsymbol{r}$ is the vector of modelled observables for the current vector of parameters, and $\boldsymbol{r}_o$ is the vector of measured observables, so an iteration step should perform a change of model parameters equal to

$$\Delta \boldsymbol{p} = \boldsymbol{V} (\boldsymbol{S}^2 + \boldsymbol{I})^{-1} [\boldsymbol{S} \boldsymbol{U}^{\mathrm{T}} (\boldsymbol{r}_o - \boldsymbol{r}) + \boldsymbol{V}^{\mathrm{T}} (\boldsymbol{p}_a - \boldsymbol{p})] . \tag{4}$$

This updating rule forms an effective cure for the ill-posedness problem, since also for singular values equal to zero a stable solution is found. In that case the corresponding linear transformations of surface variables will automatically be equated to the *a priori* ones.

4. RESULTS

Before showing some examples of successful retrievals of many parameters by the Bayesian model inversion algorithm from simulated hyperspectral-multidirectional TOA radiance data, the dimensionality of this kind of rich data (201 bands × 7 directions) is investigated by analysing the Jacobian matrix of model sensitivities for a single benchmark case. This case represents a clumped (less than 100% crown cover) and mixed vegetation canopy (green and brown leaves both present) with a crown LAI of one. This kind of object was chosen in order to give all parameters some sensitivity on remote sensing observables, so that a maximum potential dimensionality would be reached. For some settings of model variables the sensitivities to a subset of variables may go to zero. A trivial example is a canopy which only has green leaves (fraction brown leaf area zero). This gives zero sensitivity to the brown leaf optical properties. Another is the case of no clumping (homogeneous canopy with 100% crown cover). In this case the tree shape factor becomes irrelevant.
For the case of a hyperspectral-multidirectional mission the results are given by the red line (box symbol) in Fig. 7, which shows the retrieval error variances of the linearly transformed surface variables obtained from the SVD of the Jacobian matrix, relative to those valid for the case of guessing them from the *a priori* information. Here the *a priori* information was assumed to consist of the assumption that each variable was equal to its

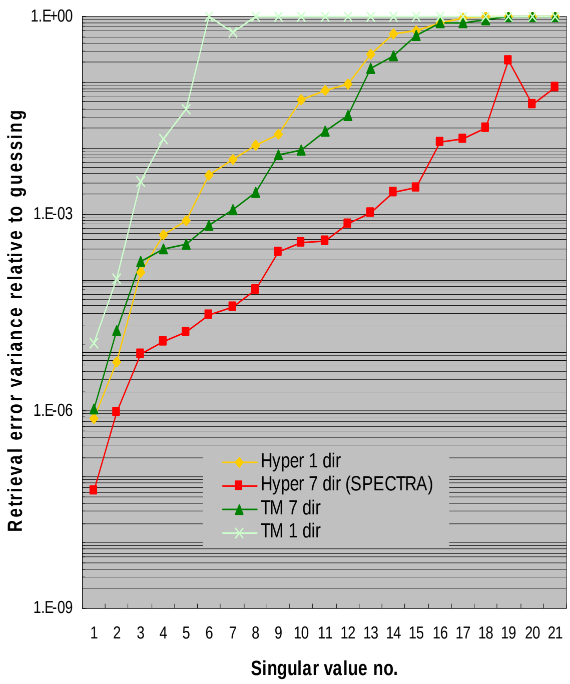

Fig. 7 Dimensionality plot for four different mission concepts.

value at the centre of its interval. The *a priori* variances (uncertainties) were expressed by assuming for each variable a uniform distribution over their entire plausible range. Dimensionality may be expressed by counting the number of linearly independent transformed variables which can be retrieved with an error variance clearly less than the error variance corresponding to guessing. For double-normalised data, i.e. the variables are normalised by their standard deviations and the observables by their noise levels (also expressed as standard deviations), the relative error variance of a linearly independent transformed variable becomes equal to $1/(S^2+1)$, where S is the associate singular value. Calling a transformed variable retrievable if its error variance is more than ten times smaller than the one corresponding to randomly guessing, in Fig. 7 one may observe that for the hyperspectral-multidirectional mission only one singular value is found that is associated with a transformed variable that should be considered non-retrievable. The total number of singular values is 21 here, since the four Hapke soil BRDF parameters were not varied, and in this case two parameters were added to include some uncertainty due to the atmospheric adjacency effect, the fractions of dense vegetation and bare soil in the neighbourhood of the target pixel. So one may conclude that the dimensionality of the data retrievable from hyperspectral-multidirectional observations in this case is 20. For a multispectral mission (6 Landsat bands) with 7 directions the dimensionality is 12, and the same number is found for a single view nadir-looking hyperspectral mission. For a nadir-looking multispectral (Landsat bands) mission the dimensionality defined in this way is 5. Note that for this simulation the noise level was assumed to be given by 0.1 $Wm^{-2}\mu m^{-1}sr^{-1}$, uniformly distributed over the spectrum. For a tenfold higher noise level (1 $Wm^{-2}\mu m^{-1}sr^{-1}$) the dimensionalities obtained for the same mission types are 13, 7, 5 and 2, respectively. Here the multispectral multidirectional mission scores higher than a single-view hyperspectral mission. One may also conclude that a good signal-to-noise ratio is essential for the successful retrieval of many geobiophysical surface variables, as well as an advanced hyperspectral-multidirectional mission concept.

In order to test the performance of a Bayesian model inversion algorithm based on the regularised Newton optimisation method as outlined in section 3, 10000 combinations of randomly selected surface variables were generated, the corresponding TOA hyperspectral-multidirectional radiance data were computed by forward modelling, and next these were provided to the retrieval algorithm. In this numerical experiment all 23 variables of the SLC model were allowed to vary and these should all be retrieved as well. However, the dry soil reflectance spectrum and the atmospheric properties were assumed to be known in this case. The result was that in 99% of the cases the correct solution was found, and that the number of Newton iterations was mostly less than twenty. Computation time on a common PC is quite considerable, since one iteration requires the computation of the Jacobian matrix, which involves 201 × 7 × 24 model simulations, so that a complete sequence of iterations for a single optimisation may well take 5 to 10 seconds. However, one could imagine that the full optimisation would be applied only to field-averaged data or certain representative pixels, while look-up table methods might be applied elsewhere to capture intra-field variations (Verhoef & Bach, 2003). In that case most variables are set to fixed values, and only a few (up to four) are allowed to vary.

An example illustrating the functioning of the algorithm for the retrieval of the brown leaf optical properties is shown in Fig. 8. This example was chosen to illustrate in particular that for some parameters the final solution deviates from the correct one because of the bias created by the *a priori* solution. The modelled TOA radiances are only weakly sensitive to some object parameters, and for parameters for which this is the case the final solution will give relatively much weight to the *a priori* values. This is illustrated in Fig. 8, where for brown leaves after 11 iterations still considerable deviations are found for the biochemical components chlorophyll, water, and brown pigment.

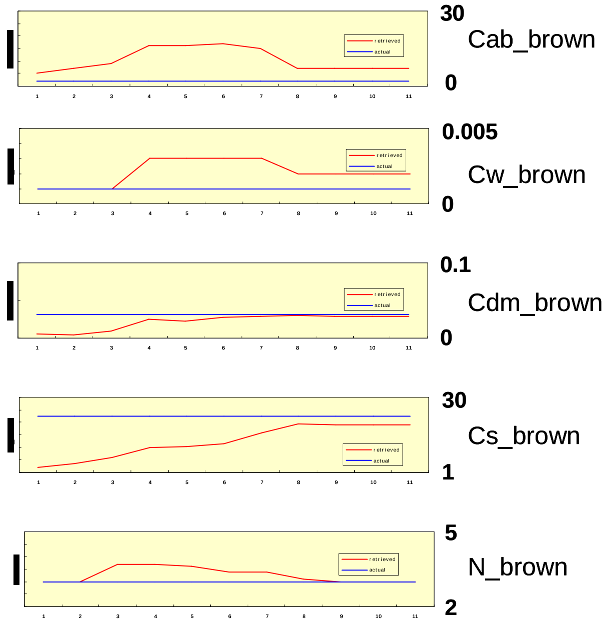

Fig. 8 Retrieval of brown leaf optical properties. Flat blue lines show the correct values to be retrieved, red lines the successive trials during 11 iteration steps. The right axis shows the names and plausible ranges of the respective parameters.

How, for the same Bayesian model inversion experiment, the most important canopy structural parameters were retrieved is shown in Fig. 9. Especially the LAI and the average leaf slope (LIDFa) are found soon. The more or less correct hot spot parameter is found only after several iterations. This is caused by the fact that the simulated observations were not in the principal plane (minimum relative azimuth about 25 deg), so sensitivity to the hot spot parameter is only small.

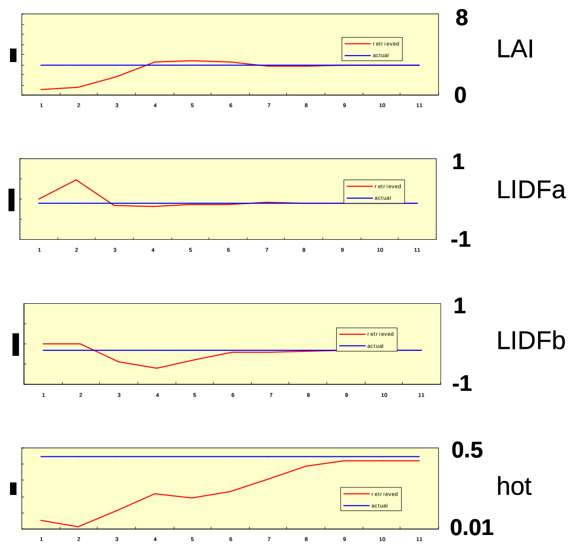

Fig. 9 Retrieval of some canopy structure parameters.

Although it was found that Bayesian retrieval of parameters by regularising Newton's method with *a priori* information, it should be noted that some extra measures were still necessary to improve the stability of the algorithm, the most important one being the initial search for a suitable starting LAI. For this, a small look-up table is constructed which contains the model results for eight LAI values and default values for the other parameters. This look-up table is used to find the best starting LAI.

Another refinement was introduced to compensate for the non-linearity of the model's response to LAI. An exponentially transformed LAI was used to improve linearity, which has a favourable effect on convergence speed.

5. CONCLUSIONS

A concept of remote sensing data assimilation has been presented which retrieves land surface information from new earth observation data by comparing this to forward-modelled existing *a priori* information, and applying a feedback loop on the level of complete (series of) images. In this mechanism the balancing of old and new information plays a key role, and for this a Bayesian approach based on the uncertainties of remote sensing observables and the *a priori* surface variables appears very attractive. The concept can be used for updating of surface variables from several sensors on board of several earth observation missions, thus bridging the gaps between sensor properties and improving the continuity and consistency of land surface products. This concept has been prototyped on the information-rich simulated data that one might expect from a hypothetical hyperspectral-multidirectional earth observation mission.

The information content of hyperspectral multidirectional radiance observations from space has been investigated by means of singular value decomposition (SVD) of the Jacobian matrix, which expresses the coupled model's sensitivity to changes in surface variables. In a Bayesian context, information content can be defined by the ability to retrieve surface variables relative to the *a priori* uncertainty of these variables. This approach allows to establish the dimensionality of the data as a function of instrument performances (noise levels) for several earth observation mission concepts. Furthermore, it can be used as the basis for an efficient and stable parameter retrieval algorithm, employing Newton's method with incorporation of the *a priori* uncertainties.

Hyperspectral multidirectional earth observation data are a very rich source of information about soil-vegetation objects. In spite of the huge number of observables ($7 \times 201 = 1407$ in our example), dimensionality appears to remain limited to 13 respectively 20, depending on the assumed noise levels. This seems to indicate that a considerable data compression by a factor of 70 should be possible. It was also found that the dimensionalities of single-view hyperspectral data and multiple view multispectral data (6 bands) are comparable and amount to about 6 and 12 for high and low noise levels, respectively. A single-view multispectral mission produces dimensionalities of 2 and 5 at the same noise levels.

A priori information was incorporated successfully to remove ill-posedness and to retrieve all 23 parameters of the SLC model by a Bayesian model inversion approach.

For randomly generated model cases, in 99% of the cases a correct solution was found in less than 20 iteration steps. In the remaining cases no convergence was achieved.

6. REFERENCES

Bach, H., Mauser, W. (1994). Modeling and model verification of the spectral reflectance of soils under varying moisture conditions. *Proc. IGARSS'94.* Vol. 4. pp. 2354-2356, Pasadena: IEEE.

Berk, A., Anderson, G.P., Acharya, P.K., Chetwynd, J.H., Bernstein, L.S., Shettle, E.P., Matthew, M.W., Adler-Golden, S.M. (2000). MODTRAN4 Version 2 USER'S MANUAL, Air Force Research Laboratory. Space Vehicles Directorate, Air Force Materiel Command, Hanscom AFB, MA 01731-3010, 97 pp.

Hapke, B.W. (1981). Bi-directional reflectance spectroscopy 1. Theory, J. Geophys. Res. 86: 3039-3054.

Jacquemoud, S., Baret, F. (1990). PROSPECT: A model of leaf optical properties spectra. *Rem. Sens. of Env.* 34: 75-91.

Knyazikhin, Y., Martonchik, J.V., Myneni, R.B., Diner, D.J., Running, S.W. (1998). Synergistic algorithm for estimating vegetation canopy leaf area index and fraction of absorbed photosynthetically active radiation from MODIS and MISR data, *Journal of Geophysical Research*, Vol. 103, No. D24, pp. 32,257-32,276.

Verhoef, W., Bach, H. (2003). Simulation of hyperspectral and directional radiance images using coupled biophysical and atmospheric radiative transfer models. *Rem. Sens. of Env.* 87: 23-41.

Verhoef, W., Bach, H. (2007). Coupled soil-leaf-canopy and atmosphere radiative transfer modeling to simulate hyperspectral multi-angular surface reflectance and TOA radiance data. *Rem. Sens. of Env.*, 17 pp., in press.

7. ACKOWLEDGEMENTS

The work reported here was supported by NLR's own research programme, as well as by the European Space Agency (ESTEC Contract No. 17169/03/NL/GS, "Impact of the Accuracy of Variables Retrievable from Multi-angular, Hyperspectral Observations on Ecosystem Model Performance", led by Massimo Menenti, with Michael Rast as the ESA study manager).